

Divide-and-Conquer Method for L_1 Norm Matrix Factorization in the Presence of Outliers and Missing Data

Deyu Meng and Zongben Xu

Abstract

The low-rank matrix factorization as a L_1 norm minimization problem has recently attracted much attention due to its intrinsic robustness to the presence of outliers and missing data. In this paper, we propose a new method, called the divide-and-conquer method, for solving this problem. The main idea is to break the original problem into a series of smallest possible sub-problems, each involving only unique scalar parameter. Each of these sub-problems is proved to be convex and has closed-form solution. By recursively optimizing these small problems in an analytical way, efficient algorithm, entirely avoiding the time-consuming numerical optimization as an inner loop, for solving the original problem can naturally be constructed. The computational complexity of the proposed algorithm is approximately linear in both data size and dimensionality, making it possible to handle large-scale L_1 norm matrix factorization problems. The algorithm is also theoretically proved to be convergent. Based on a series of experiment results, it is substantiated that our method always achieves better results than the current state-of-the-art methods on L_1 matrix factorization calculation in both computational time and accuracy, especially on large-scale applications such as face recognition and structure from motion.

Index Terms

Low-rank matrix factorization, robustness, divide-and-conquer, face recognition, structure from motion.

Deyu Meng and Zongben Xu are with the Institute for Information and System Sciences, Xi'an Jiaotong University, Xi'an, 710049, P.R.

China. (dymeng@xjtu.edu.cn)

I. INTRODUCTION

In recent years, the problem of low-rank matrix factorization (LRMF) has attracted much attention due to its wide range of applications in computer vision and pattern recognition, such as structure from motion [15], face recognition [16], shape from varying illumination [8], and object tracking [1]. Representing the measurements or the observation data as a $d \times n$ matrix $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, whose columns \mathbf{x}_i s correspond to the d -dimensional input measurements and n is the number of input items, the aim of the LRMF can be mathematically described as solving the following optimization problem:

$$\min_{U,V} \|X - UV^T\|, \quad (1)$$

where $U = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k) \in R^{d \times k}$, $V = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k) \in R^{n \times k}$ and $k < d, n$. To deal with the real LRMF problems in the presence of missing data, the optimization (1) is also reformulated as

$$\min_{U,V} \|W \odot (X - UV^T)\|, \quad (2)$$

where \odot denotes the component-wise multiplication (i.e., the Hadamard product), and the element w_{ij} of the denotation matrix $W \in R^{d \times n}$ is 1 if the corresponding element of X is known, and 0 otherwise [2]. Here $\|\cdot\|$ is some form of the matrix norm.

The global minimum of the optimization problem (1) with L_2 matrix norm (i.e., the Frobenius norm) can easily be solved by the well known singular value decomposition (SVD, [7]) method. To handle missing data, some methods, such as the Wiberg algorithm [13] and the weighted low-rank approximation method (WLRA, [14]), have further been proposed to solve the optimization (2) with L_2 matrix norm. The performance of these techniques, however, is sensitive to the presence of outliers or noises, which often happen in real measurements, because the influence of outliers or noises with a large norm tends to be considerably exaggerated by the use of the L_2 norm [3], [11]. To alleviate this robustness problem, the often used approach is to replace the L_2 matrix norm with the L_1 norm in the objective functions of (1) and (2) [5], [4], [9], [11]. The models are then expressed in the following forms:

$$\min_{U,V} \|X - UV^T\|_{L_1} \quad (3)$$

and

$$\min_{U,V} \|W \odot (X - UV^T)\|_{L_1}. \quad (4)$$

Unfortunately, it turns out that replacing the L_2 norm with L_1 norm in the optimizations makes the problem significantly more difficult [5]. First, both (3) and (4) are non-convex problems, so their global optimality are in general difficult to obtain. Second, both optimizations are non-smooth problems conducted by the L_1 matrix norm, so it is hard to attain an easy closed-form iteration formula to efficiently approximate their solutions by standard optimization tools. Third, in real applications, both optimizations can also be very computationally demanding problems to solve, which always limit their availability in large-scale practice.

In this paper, by employing an important algorithm design paradigm, namely divide and conquer, we formulate efficient algorithms against the optimization models (3) and (4), respectively. The core idea underlying the new algorithms is to break the original optimizations into a series of smallest possible problems and recursively solve them. Each of these small problems is convex and has closed-form solution, which enables the new algorithms to avoid using a time-consuming numerical optimization as an inner loop. The proposed algorithms are thus easy to implement. Especially, it is theoretically evaluated that the computational speeds of the proposed algorithms are approximately linear in both data size and dimensionality, which allows them to handle large-scale L_1 norm matrix factorization problems. The efficiency and robustness of the proposed algorithms have also been substantiated by a series of experiments implemented on synthetic and real data.

Throughout the paper, we denote matrices, vectors, and scalars by the upper-case letters, lower case bold-faced letters, and lower-case non-bold-faced letters, respectively.

II. RELATED WORK

Various approaches have recently been proposed to deal with the optimizations (3) and (4) to achieve robust low-rank matrix factorization results. For the L_1 norm model (3), the iteratively re-weighted least-squares approach introduced by Torre and Black is one of the first attempts [3]. Its main idea is to iteratively

assign a weight to each element in the measurements. The method, however, is generally very sensitive to initialization ([9]). Instead of the L_1 matrix norm, Ding et al. utilized the rotational invariant R_1 norm, as defined by $\|X\|_{R_1} = \sum_{i=1}^n (\sum_{j=1}^d x_{ji}^2)^{1/2}$, for the objective function of (3) (ICML06, [4]). Like the L_1 norm, the R_1 norm so defined is also capable of softening the contributions from outliers. By substituting the maximization of the L_1 dispersion of data, $\|U^T X\|_{L_1}$, for the minimization of the original L_1 objective, $\|X - UV^T\|_{L_1}$, Kwak presented another approach for the problem (PAMI08, [11]). The method is also able to suppress the negative effects of outliers to a certain extent. The most predominance of this method is its fast computational speed, which is linear in both measurement size and dimensionality.

Two methods represent the current state of the art of solving the model (4). The first method is presented by Ke and Kanade, who formulated the robust L_1 norm matrix factorization objective as alternative convex programs (CVPR05, [9]). The programs can then be efficiently solved by linear or quadratic programming. The second method is designed by Eriksson and Hengel, which represents a generalization of the traditional Wiberg algorithm (CVPR10, [5]). The method has been empirically proved to perform well on some synthetic and real world problems, such as the structure from motion (SFM) applications. It should be noted that both methods can also be employed to solve (3) by setting all elements of the missing data denotation matrix W to be 1s.

III. DIVIDE-AND-CONQUER METHOD FOR ROBUST MATRIX FACTORIZATION

Unlike the previous methods for robust matrix factorization, the proposed divide-and-conquer (D&C in brief) method chooses to solve the smallest possible sub-problems of (3) and (4) at every step (involving only one scalar parameter of U or V). The advantage of the new method lies in the fact that each small sub-problem so attained can be solved analytically. Thus, complicated numerical optimization techniques are entirely avoided, and the overall problem can thus be efficiently solved. We introduce our method and its theoretical fundament as follows.

A. Breaking the model into smallest sub-problems

We first consider the optimization model (3). Since the k -rank matrix UV^T can be partitioned into the sum of k 1-rank matrices, i.e., $UV^T = \sum_{i=1}^k \mathbf{u}_i \mathbf{v}_i^T$, (3) can thus be equivalently reformulated as

$$\min_{\{\mathbf{u}_j, \mathbf{v}_j\}_{j=1}^k} \left\| X - \sum_{j=1}^k \mathbf{u}_j \mathbf{v}_j^T \right\|_{L_1}. \quad (5)$$

The original k -rank matrix factorization problem can then be decomposed into a series of recursive 1-rank sub-problems:

$$\min_{\mathbf{u}_i, \mathbf{v}_i} \|E_i - \mathbf{u}_i \mathbf{v}_i^T\|_{L_1}, \quad (6)$$

where $E_i = (\mathbf{e}_1^i, \mathbf{e}_2^i, \dots, \mathbf{e}_n^i) = X - \sum_{j \neq i} \mathbf{u}_j \mathbf{v}_j^T$. We can naturally approximate the solution of (5) by sequentially solving (6) with respect to $(\mathbf{u}_i, \mathbf{v}_i)$ for $i = 1, 2, \dots, k$, with all other $(\mathbf{u}_j, \mathbf{v}_j)$ s ($j \neq i$) fixed.

Solving (6) can further be simplified to alternatively optimizing \mathbf{u}_i or \mathbf{v}_i while letting the other fixed. Since \mathbf{u}_i and \mathbf{v}_i can be solved in a completely symmetrical way (in the sense that $\|E - \mathbf{u} \mathbf{v}^T\|_{L_1} = \|E^T - \mathbf{v} \mathbf{u}^T\|_{L_1}$), we only need to consider how to efficiently solve

$$\min_{\mathbf{v}_i} \|E_i - \mathbf{u}_i \mathbf{v}_i^T\|_{L_1}. \quad (7)$$

By reformulating (7) to its decoupling form:

$$\min_{\mathbf{v}_i = (v_{i1}, v_{i2}, \dots, v_{in})^T} \sum_{j=1}^n \|\mathbf{e}_j^i - \mathbf{u}_i v_{ij}\|_{L_1}, \quad (8)$$

where v_{ij} is the j -th element of the vector \mathbf{v}_i , the problem can then be further divided into n small sub-optimizations with the following expression (for $j = 1, 2, \dots, n$):

$$\min_{v_{ij}} \|\mathbf{e}_j^i - \mathbf{u}_i v_{ij}\|_{L_1}. \quad (9)$$

From (5) to (9), we have broken the original large optimization (3), with respect to U and V , into a series of smallest possible optimization problems, each with respect to only one scalar parameter of U or V . By utilizing the similar strategy, it is also easy to decompose the large optimization (4) into a series of small optimizations, expressed as:

$$\min_{v_{ij}} \|\mathbf{w}_j \odot (\mathbf{e}_j^i - \mathbf{u}_i v_{ij})\|_{L_1}, \quad (10)$$

where \mathbf{w}_j is the j -th column of the denotation matrix W .

It is very fortunate that both small optimizations (9) and (10) are not only convex, but also have closed form solutions. This implies that it is possible to construct fast algorithms for (3) and (4), as introduced in the following discussion.

B. The closed form solutions of (9) and (10)

We first formalize (9) as:

$$\min_v f_{\mathbf{e}, \mathbf{u}}(v) = \|\mathbf{e} - \mathbf{u}v\|_{L_1} \quad (11)$$

where both \mathbf{e} and \mathbf{u} are d -dimensional vectors, and denote their i -th elements as e_i and u_i , respectively. The following theorem shows the convexity of this optimization problem (the proofs of all involved theorems are moved to the supplementary material due to the page limitation).

Theorem 1: $f_{\mathbf{e}, \mathbf{u}}(v)$ as defined in (11) is a convex function with respect to v .

Theorem 1 implies that it is hopeful to find the global optimum of (11). We first clarify the case when all elements of \mathbf{u} are positive in the following lemma.

Lemma 1: For (11), assuming each element u_i of \mathbf{u} is positive ($u_i > 0, i = 1, 2, \dots, d$), denote

- the label set $\mathbf{L}_{\mathbf{e}, \mathbf{u}} = (l_1^{(\mathbf{e}, \mathbf{u})}, l_2^{(\mathbf{e}, \mathbf{u})}, \dots, l_d^{(\mathbf{e}, \mathbf{u})})$: the permutation of $(1, 2, \dots, d)$ based on the ascending order of $(\frac{e_1}{u_1}, \frac{e_2}{u_2}, \dots, \frac{e_d}{u_d})$;
- the sequence $\mathbf{\Gamma}_{\mathbf{e}, \mathbf{u}} = (a_0, a_1, \dots, a_d)$: $a_0 = -\sum_{j=1}^d u_{l_j^{(\mathbf{e}, \mathbf{u})}}$, $a_d = \sum_{j=1}^d u_{l_j^{(\mathbf{e}, \mathbf{u})}}$, $a_i = \sum_{j=1}^i u_{l_j^{(\mathbf{e}, \mathbf{u})}} - \sum_{j=i+1}^d u_{l_j^{(\mathbf{e}, \mathbf{u})}}$, $i = 1, 2, \dots, d-1$;
- the label $i_{\mathbf{e}, \mathbf{u}}$: the label of the first non-negative element of $\mathbf{\Gamma}_{\mathbf{e}, \mathbf{u}}$;

and the following closed form expression provides a global optimum of (11):

$$P(\mathbf{e}, \mathbf{u}) := \frac{e_{i^*}}{u_{i^*}}, \quad i^* = l_{i_{\mathbf{e}, \mathbf{u}}}^{(\mathbf{e}, \mathbf{u})}. \quad (12)$$

It is easy to deduce that $\mathbf{\Gamma}_{\mathbf{e}, \mathbf{u}}$ is a monotonically increasing sequence, and $a_0 < 0$, $a_d > 0$. Thus, the label $i_{\mathbf{e}, \mathbf{u}}$ can be uniquely found from the sequence.

Algorithm 1: D&C algorithm for solving $\min_{U,V} \|X - UV^T\|_{L_1}$

Given: $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) \in R^{d \times n}$

Execute:

1. Randomly initialize $U^{(0)} = (\mathbf{u}_1^{(0)}, \mathbf{u}_2^{(0)}, \dots, \mathbf{u}_k^{(0)}) \in R^{d \times k}$, $V^{(0)} = (\mathbf{v}_1^{(0)}, \mathbf{v}_2^{(0)}, \dots, \mathbf{v}_k^{(0)}) \in R^{n \times k}$;
 - For** $t = 1, 2, \dots$, *convergence*
 2. Let $\mathbf{u}_i^{(t)} = \mathbf{u}_i^{(t-1)}$, $\mathbf{v}_i^{(t)} = \mathbf{v}_i^{(t-1)}$, $i = 1, 2, \dots, k$.
 - For** $i = 1, 2, \dots, k$
 3. Compute $E_i = X - \sum_{j=1, j \neq i}^k \mathbf{u}_j^{(t)} \mathbf{v}_j^{(t)T}$; denote the column and row vectors of E_i as $(\mathbf{e}_1^i, \mathbf{e}_2^i, \dots, \mathbf{e}_n^i)$ and $(\tilde{\mathbf{e}}_1^i, \tilde{\mathbf{e}}_2^i, \dots, \tilde{\mathbf{e}}_d^i)$, respectively.
 4. Let $v_{ij}^{(t)} = Q(\mathbf{e}_j^i, \mathbf{u}_i^{(t)})$ for $j = 1, 2, \dots, n$ based on (13), and then update $\mathbf{v}_i^{(t)} = (v_{i1}^{(t)}, v_{i2}^{(t)}, \dots, v_{in}^{(t)})^T$.
 5. Let $u_{ij}^{(t)} = Q(\tilde{\mathbf{e}}_j^i, \mathbf{v}_i^{(t)})$ for $j = 1, 2, \dots, d$ based on (13), and then update $\mathbf{u}_i^{(t)} = (u_{i1}^{(t)}, u_{i2}^{(t)}, \dots, u_{id}^{(t)})^T$.
 - End For**
 6. Update $U^{(t)} = (\mathbf{u}_1^{(t)}, \mathbf{u}_2^{(t)}, \dots, \mathbf{u}_k^{(t)})$, $V^{(t)} = (\mathbf{v}_1^{(t)}, \mathbf{v}_2^{(t)}, \dots, \mathbf{v}_k^{(t)})$.
 - End For**
-

The above theorem gives a closed form solution for (11) under positive vector \mathbf{u} . Next theorem further gives the solution of (11) in general cases.

Theorem 2: For (11), denote

- the label set $\mathbf{I}_{\mathbf{u}} = (i_1, i_2, \dots, i_{\hat{d}})(\hat{d} \leq d)$: the labels of the nonzero elements of \mathbf{u} ;
- $\Upsilon_{\mathbf{u}} = (u_{i_1}, u_{i_2}, \dots, u_{i_{\hat{d}}})^T$, $\Psi_{\mathbf{e}, \mathbf{u}} = (e_{i_1}, e_{i_2}, \dots, e_{i_{\hat{d}}})^T$;
- $\tilde{\Upsilon}_{\mathbf{u}} = \text{sign}(\Upsilon_{\mathbf{u}}) \odot \Upsilon_{\mathbf{u}}$, $\tilde{\Psi}_{\mathbf{e}, \mathbf{u}} = \text{sign}(\Upsilon_{\mathbf{u}}) \odot \Psi_{\mathbf{e}, \mathbf{u}}$, where $\text{sign}(\cdot)$ is the signum function;

and the following closed form expression provides a global optimum of (11):

$$Q(\mathbf{e}, \mathbf{u}) := P(\tilde{\Psi}_{\mathbf{e}, \mathbf{u}}, \tilde{\Upsilon}_{\mathbf{u}}), \quad (13)$$

where the function $P(\cdot, \cdot)$ is defined as (12).

We then consider (10). First formalize it as

$$\min_v f_{\mathbf{w}, \mathbf{e}, \mathbf{u}}(v) = \|\mathbf{w} \odot (\mathbf{e} - \mathbf{u}v)\|_{L_1}, \quad (14)$$

Algorithm 2: D&C algorithm for solving $\min_{U,V} \|W \odot (X - UV^T)\|_{L_1}$

4. Let $v_{ij}^{(t)} = Q(\mathbf{w}_j \odot \mathbf{e}_j^i, \mathbf{w}_j \odot \mathbf{u}_i^{(t)})$ for $j = 1, 2, \dots, n$ based on (15), and then update $\mathbf{v}_i^{(t)} = (v_{i1}^{(t)}, v_{i2}^{(t)}, \dots, v_{in}^{(t)})^T$.
 5. Let $u_{ij}^{(t)} = Q(\tilde{\mathbf{w}}_j \odot \tilde{\mathbf{e}}_j^i, \tilde{\mathbf{w}}_j \odot \mathbf{v}_i^{(t)})$ for $j = 1, 2, \dots, d$ based on (15), and then update $\mathbf{u}_i^{(t)} = (u_{i1}^{(t)}, u_{i2}^{(t)}, \dots, u_{id}^{(t)})^T$.
-

where \mathbf{w} , \mathbf{e} , \mathbf{u} are all d -dimensional vectors. Since

$$\|\mathbf{w} \odot (\mathbf{e} - \mathbf{u}v)\|_{L_1} = \|(\mathbf{w} \odot \mathbf{e}) - (\mathbf{w} \odot \mathbf{u})v\|_{L_1},$$

(14) can then be seen as a special case of (11) in the sense that

$$f_{\mathbf{w}, \mathbf{e}, \mathbf{u}}(v) = f_{\mathbf{w} \odot \mathbf{e}, \mathbf{w} \odot \mathbf{u}}(v).$$

It thus holds the following theorem based on Theorems 1 and 2.

Theorem 3: (14) is a convex optimization problem with closed form solution

$$Q(\mathbf{w} \odot \mathbf{e}, \mathbf{w} \odot \mathbf{u}), \tag{15}$$

where $Q(\cdot, \cdot)$ is defined as (13).

By virtue of the closed form solutions for the small optimization problems (9)/(11) and (10)/(14) given by Theorems 2 and 3, respectively, we can now construct fast algorithms for solving the original large robust matrix factorization problems (3) and (4).

C. Fast algorithms for robust matrix factorization

We first consider the D&C algorithm for the optimization model (3). The main idea of our algorithm is to sequentially update each element of U and V . In specific, the algorithm iteratively updates each element of \mathbf{u}_i and \mathbf{v}_i for $i = 1, 2, \dots, k$, with other \mathbf{u}_j s and \mathbf{v}_j s ($j \neq i$) fixed, in the following way:

- Update each element v_{ij} ($j = 1, 2, \dots, n$) of \mathbf{v}_i under fixed \mathbf{u}_i : the optimal value v_{ij} is attained through the following closed form expression based on Theorem 2

$$v_{ij} = Q(\mathbf{e}_j^i, \mathbf{u}_i) = \arg \min_{v_{ij}} \|\mathbf{e}_j^i - \mathbf{u}_i v_{ij}\|_{L_1},$$

where \mathbf{e}_j^i is the j -th column vector of the representation error matrix

$$E_i = X - \sum_{j \neq i} \mathbf{u}_j \mathbf{v}_j^T. \quad (16)$$

- Update each element u_{ij} ($j = 1, 2, \dots, d$) of \mathbf{u}_i under fixed \mathbf{v}_i : the optimal value of u_{ij} is achieved through

$$u_{ij} = Q(\tilde{\mathbf{e}}_j^i, \mathbf{v}_i) = \arg \min_{u_{ij}} \|\tilde{\mathbf{e}}_j^i - \mathbf{v}_i u_{ij}\|_{L_1}$$

based on Theorem 2, where $\tilde{\mathbf{e}}_j^i$ denotes the j -th row vector of E_i .

Through implementing the above iterations from $i = 1$ to k , the low-rank factorized matrices U and V can then be recursively updated until the termination condition is satisfied. We embed the aforementioned D&C technique into Algorithm 1.

The D&C algorithm for solving (4) is very similar to Algorithm 1, the only difference is the updating of each element of \mathbf{u}_i and \mathbf{v}_i , which is summarized as follows:

- Update each element v_{ij} ($j = 1, 2, \dots, n$) of \mathbf{v}_i under fixed \mathbf{u}_i : the optimal value v_{ij} is solved through

$$\begin{aligned} v_{ij} &= Q(\mathbf{w}_j \odot \mathbf{e}_j^i, \mathbf{w}_j \odot \mathbf{u}_i) \\ &= \arg \min_{v_{ij}} \|\mathbf{w}_j \odot (\mathbf{e}_j^i - \mathbf{u}_i v_{ij})\|_{L_1} \end{aligned}$$

based on Theorem 3, where \mathbf{w}_j is the j -th column vector of W , and \mathbf{e}_j^i is the j -th column vector of the representation error matrix E_i (defined as (16)).

- Update each element u_{ij} ($j = 1, 2, \dots, d$) of \mathbf{u}_i under fixed \mathbf{v}_i : the optimal value of u_{ij} is attained by

$$\begin{aligned} u_{ij} &= Q(\tilde{\mathbf{w}}_j \odot \tilde{\mathbf{e}}_j^i, \tilde{\mathbf{w}}_j \odot \mathbf{v}_i) \\ &= \arg \min_{u_{ij}} \|\tilde{\mathbf{w}}_j \odot (\tilde{\mathbf{e}}_j^i - \mathbf{v}_i u_{ij})\|_{L_1} \end{aligned}$$

based on Theorem 3, where $\tilde{\mathbf{w}}_j$ denotes the j -th row vector of W , and $\tilde{\mathbf{e}}_j^i$ is the j -th row vector of E_i .

Since the D&C algorithm for solving (4) differs from Algorithm 1 only in steps 4 and 5, we only list these two steps of the algorithm in Algorithm 2.

The remaining issues are then on how to appropriately specify the initial U and V in step 1, and when to terminate the iterative process in steps 2-6 of the proposed algorithms. In our experiments, we just randomly initiated each element of U and V , and the proposed algorithm performed well in all experiments under such simple initialization strategy. As for the termination of the algorithms, since the objective function of (3) or (4) decreases monotonically throughout the iterative process (see details in the next section), the algorithms can reasonably be terminated when the updating extent $\|U^{(t)} - U^{(t-1)}\|$ (or $\|V^{(t)} - V^{(t-1)}\|$) is smaller than some preset small threshold, or the process has reached the pre-specified number of iterations.

D. Convergence and computational complexity

We now discuss the convergence of the proposed algorithms. It should be noted that in steps 4 and 5 of Algorithm 1/Algorithm 2, the global minimum of the objective function of (3)/(4) with respect to \mathbf{v}_i and \mathbf{u}_i (with other \mathbf{v}_j s and \mathbf{u}_j s fixed) is analytically obtained based on Theorem 2/Theorem 3, respectively, and thus in each of the iteration steps of the algorithm, the objective function of the problem is monotonically decreasing. Since it is evident that both objective functions of (3) and (4) are lower bounded (≥ 0), the algorithm is guaranteed to be convergent.

The computational complexity of Algorithm 1/Algorithm 2 is essentially determined by the iterations between steps 4 and 5, i.e., the calculation of the closed form solutions of \mathbf{v}_i and \mathbf{u}_i . To compute the global optimum for each element v_{ij} of \mathbf{v}_i ($j = 1, 2, \dots, n$) in step 4 of the algorithm, the closed form expression $P(\cdot, \cdot)$, as defined in Lemma 1, is utilized, which costs $O(d \log d)$ computation to obtain the label set $\mathbf{L}_{\mathbf{e}, \mathbf{u}}$ by applying the well-known heap sorting algorithm [10], and costs at most $O(d)$ computation to seek the label of the first nonzero element $i_{\mathbf{e}, \mathbf{u}}$ from the sequence $\Gamma_{\mathbf{e}, \mathbf{u}}$. Altogether, calculating the global optimum for each v_{ij} needs around $O(d \log d)$ computational time. Updating the entire \mathbf{v}_i thus requires about $O(nd \log d)$ computational cost. It can similarly be deduced that updating \mathbf{u}_i in step 5 needs

around $O(nd \log n)$ computational cost. The entire computational complexity of Algorithm 1/Algorithm 2 is thus about $O(k(nd \log d + nd \log n)) \times T$, where T is the number of iterations for convergence. That is, the computational speeds of the proposed algorithms are approximately linear in both the size and dimensionality of the input measurements X , as well as its intrinsic rank k . Such computational complexity makes the use of the proposed algorithms possible in large-scale L_1 norm matrix factorization problems, as demonstrated in the following experiments.

IV. EXPERIMENTS

To evaluate the performance of the proposed D&C algorithms on robust matrix factorization, it was applied to various synthetic and real problems with outliers and missing data. The results are summarized in the following discussion. All programs were implemented under the Matlab 7.0 platform. The implementation environment was the personal computer with Intel Core(TM)2 Quad Q9300@2.50 G (CPU), 3.25GB (memory), and Windows XP (OS).

A. Experiments on data with outliers

Three series of experiments were designed to evaluate the performance of the proposed Algorithm 1 on data with intrinsic outliers (for solving the optimization model (3)). The details are listed as follows:

Small synthetic experiment E1: Containing 100 synthetic 30×30 matrices, each with intrinsic rank 3. Each matrix was first generated as a product UV^T , where U and V are independent 30×3 matrices, whose elements are i.i.d. Gaussian random variables with zero mean and unit variance; and then 10% elements of the matrix were randomly picked up and transformed into outliers by randomly assigning them values in the range of $[-40, 40]$.

Large synthetic experiments E2, E3, E4: Containing 3 synthetic 7000×7000 matrices, each with intrinsic rank 3. Each was first generated from the product UV^T , where U and V are independent 7000×3 matrices, whose elements are i.i.d. Gaussian random variables with zero mean and unit variance, and then different




Fig. 1. From top to bottom: Typical Yale B face images, the corresponding images corrupted with 20% outliers, and the faces reconstructed by the SVD, the PAMI08, and the proposed D&C methods, respectively.

extents of outliers were assigned to randomly selected 10% elements of the original matrices, with ranges $[-40,40]$, $[-400,400]$, and $[-4000,4000]$, in E2, E3, and E4, respectively.

Face recognition experiments E5,E6: The input data are composed by 256 face images, each with pixel size 192×168 , i.e., the matrix is of the size 32256×256 . The images were first extracted from subsets 1-4 of the Extended Yale B database [6], and then were corrupted with 20% and 50% of dead pixels with either maximum or minimum intensity values in E5 and E6, respectively. Typical images are depicted in Figures 1 and 2.

In each of these experiments, the original un-corrupted matrix, denoted as X , is saved as ground truth for comparison purpose.

For comparison, 5 of the current methods for low-rank matrix factorization, including SVD, ICML06




Fig. 2. From top to bottom: Typical Yale B face images, the corresponding images corrupted with 50% outliers, and the faces reconstructed by the SVD, the PAMI08, and the proposed D&C methods, respectively.

[4], PAMI08 [11], CVPR05 [9], and CVPR10 [5], have also been utilized. Except SVD, which need not be initialized, and ICML06, which requires the SVD initialization [4], all of the utilized methods employed the similar initialization for each involved experiment. The rank k was preset as 3 in E1-E4 and 20 in E5 and E6 for all methods. The performance comparison of these methods is shown in Table I (that of E1 is shown as the average result over 100 experiments). In the table, / means that the corresponding method on the experiment could not be completed in reasonable time. For easy observation, Figures 1 and 2 demonstrate some of the original and reconstructed images in E5 and E6, respectively. The reconstruction is implemented by the product of $\tilde{U}\tilde{V}^T$, where the low-rank matrices \tilde{U} and \tilde{V} are the outputs of the corresponding matrix factorization method.

The advantage of the proposed Algorithm 1, as compared with other utilized methods, can evidently be

	Computational time (s)						Accuracy: $\frac{\ X - \tilde{U}\tilde{V}^T\ _F}{\ X\ _F}$ ([12])					
E1	0.0045	0.022	0.0014	0.835	411.0	0.048	7.67	7.64	5.95	0.0693	0.0221	3.57×10^{-4}
E2	137.31	/	139.06	/	/	7612	0.025	/	0.030	/	/	7.08×10^{-16}
E3	146.13	/	159.60	/	/	6953	3.91	/	3.24	/	/	2.04×10^{-12}
E4	119.85	/	189.54	/	/	7279	293.4	/	222.9	/	/	4.39×10^{-14}
E5	33.41	/	92.64	/	/	7335	0.124	/	0.117	/	/	0.0312
E6	59.90	/	234.78	/	/	7275	0.384	/	0.338	/	/	0.0959

TABLE I

THE PERFORMANCE COMPARISON OF THE 5 CURRENT MATRIX FACTORIZATION METHODS AND THE PROPOSED D&C METHOD IN EXPERIMENTS E1-E6. IN EACH CELL, THE VALUES FROM THE LEFT TO THE RIGHT REFER TO THE SVD, ICML06, PAMI08, CVPR05, CVPR10, AND D&C METHODS, RESPECTIVELY. THE BEST RESULT IN EACH EXPERIMENT IS HIGHLIGHTED. X DENOTES THE ORIGINAL UN-CORRUPTED MATRIX (GROUND TRUTH), AND \tilde{U} AND \tilde{V} DENOTE THE OUTPUTS OF THE CORRESPONDING MATRIX FACTORIZATION METHOD. / MEANS THAT THE CORRESPONDING METHOD ON THE EXPERIMENT COULD NOT BE COMPLETED IN REASONABLE TIME.

observed from Table I in robust matrix factorization calculation. Specifically, our method attains the highest computational accuracy in all of the involved experiments. For Yale B experiments E5 and E6, it is very interesting that some latent features underlying the original faces can be extracted from the reconstructed images, as clearly depicted in Figures 1 and 2. Even more interesting is that these reconstructions are obtained from the corrupted but not the original images. The new method is thus potentially useful for latent feature extraction from noisy measurements in real applications. Another merit of our method is that it has stable performance on different extents of outliers, which can evidently be observed in the E2-E4 results, in which the reconstructed low-rank matrix attained by our method is always extremely close to the ground truth. Although the computational speed of the proposed algorithm is slower than the SVD and PAMI08 methods in the experiments, considering that both SVD and PAMI08 are not designed against the L_1 norm matrix factorization model (3), the efficiency of the proposed method is still dominant in the methods against (3), especially in large-scale cases.

B. Experiments on data with outliers and missing components

We also implemented three series of experiments to evaluate the performance of Algorithm 2 on data with intrinsic outliers and missing components (for solving the optimization model (4)). The details are summarized as follows:

Small synthetic experiment E7: Containing 100 synthetic 20×30 matrices, each with intrinsic rank 3. Each matrix was first generated as a product UV^T , where U and V are 20×3 and 30×3 matrices, respectively, whose elements were generated from the Gaussian distribution with zero mean and unit variance. Then 5% of the elements were selected at random and designated as missing by setting the corresponding entry in the matrix W to zero. To simulate outliers, uniformly distributed noises over $[-5, 5]$ were additionally added to 10% of the elements of the matrix.

Large synthetic experiment E8: Containing one synthetic 10000×700 data matrix, with intrinsic rank 40. The matrix was first generated as a product UV^T , where U and V are 10000×40 and 700×40 matrices, randomly generated from the Gaussian distribution with zero mean and unit variance. Then 20% and 10% of the elements were selected at random and designated as missing components and outliers (randomly chosen in $[-5, 5]$), respectively.

Structure from motion experiments E9, E10, E11: The structure from motion (SFM) problem can be posed as a typical low-rank matrix approximation task [5], [9]. In this series of experiments, we employ two well known SFM data sequence, the dinosaur sequence, available at <http://www.robots.ox.ac.uk/~vgg/>, and the pingpong ball sequence, available at <http://vasc.ri.cmu.edu/idb/>, for substantiation. The entire dinosaur and pingpong sequence contain projections of 4983 and 839 points tracked over 36 and 226 frames, respectively, composing 4983×72 and 839×452 SFM matrices correspondingly. Each matrix contains more than 80% missing data due to occlusions or tracking failures. As considering robust approximation in this work, we further include outliers uniformly generated from $[-5000, 5000]$ in 10% components of two matrices to form the input data of the experiments E10 and E11, respectively¹. Since some other

¹The components of both SFM matrices are also approximately located in $[-5000, 5000]$.

	Computational time (s)					Accuracy: $\frac{\ X - \tilde{U}\tilde{V}^T\ _F}{\ X\ _F}$ (E7,E8), $\frac{\ W \odot (X - \tilde{U}\tilde{V}^T)\ _F}{\ W \odot X\ _F}$ (E9-E11)				
E7	61.50	24.02	3.53	1781	0.221	6.5864	15.0713	0.3274	0.3051	0.2626
E8	134070	/	/	/	90766	0.0130	/	/	/	4.9173×10^{-19}
E9	2.9447	127.37	24.93	132.79	10.071	0.4539	0.7749	0.0426	0.0405	0.0031
E10	202.49	/	13788	/	61.917	0.4462	/	0.3385	/	0.0765
E11	224.91	/	718.82	/	70.950	0.3498	/	0.0903	/	0.0151

TABLE II

THE PERFORMANCE COMPARISON OF THE 4 CURRENT MATRIX FACTORIZATION METHODS AND THE PROPOSED D&C METHOD IN EXPERIMENTS E7-E11. IN EACH CELL, THE VALUES FROM THE LEFT TO THE RIGHT REFER TO THE WLRA, WIBERG, CVPR05, CVPR10, AND D&C METHODS, RESPECTIVELY. THE BEST RESULT IN EACH EXPERIMENT IS HIGHLIGHTED. X DENOTES THE ORIGINAL UN-CORRUPTED MATRIX, AND \tilde{U} AND \tilde{V} DENOTE THE OUTPUTS OF THE CORRESPONDING MATRIX FACTORIZATION METHOD. / MEANS THAT THE CORRESPONDING METHOD ON THE EXPERIMENT COULD NOT BE COMPLETED IN REASONABLE TIME.

robust matrix factorization methods cannot be made available at such data scales (see Table II), we further picked up 336 points from the dinosaur sequence to form a smaller 336×72 matrix, and also added 10% outliers to it to compose the input measurements of the experiment E9.

As the experiments E1-E6, the original un-corrupted matrix, denoted as X , is saved as ground truth in each experiment for comparison purpose.

Four current low-rank matrix factorization methods were employed for comparison. They include the WLRA [14] and Wiberg methods [13], which are typical methods designed for the L_2 norm model (2), the CVPR05 [9] and CVPR10 [5] methods, which are current state-of-the-art methods for solving the L_1 norm model (4). All of the utilized methods adopted similar initialization for each of the involved experiments. The performances of these methods are compared in Table II (that of E7 is depicted as the average result over 100 experiments).

The advantage of the proposed D&C method is evident based on Table II, in terms of both computational speed and accuracy. On one hand, our algorithm always attains the most accurate reconstruction of the original data matrix by the product of the obtained low-rank matrices \tilde{U} and \tilde{V} , and on the other hand, the

computation cost of the proposed algorithm is the smallest of all employed methods in most experiments (except being the second smallest in E9). It is very impressive that the computational speed of the proposed algorithm is even faster than the WLRA and the Wiberg methods, which are constructed for L_2 norm matrix factorization model, in most cases (except slower than WLRA in E9). Considering the difficulty of solving the L_1 model due to its non-convexity and non-smoothness, the efficiency of the proposed method is more prominent.

V. CONCLUSION

In this paper we have tried a new methodology, the divide and conquer technique, for solving the L_1 norm low-rank matrix factorization problems (3) and (4). The main idea is to break the original large problems into smallest possible sub-problems, each involving only one unique scalar parameter. We have proved that these sub-problems are convex, and have closed form solutions. Inspired by this theoretical result, fast algorithms have been constructed to handle the original large problems, entirely avoiding the complicated numerical optimization for the inner loops of the iteration. In specific, we have proved that the computational complexity of the new algorithms is approximately linear in both the size and dimensionality of the input data, which enables the possible utilization of the new algorithms in large-scale L_1 norm matrix factorization problems. The convergence of the new algorithms have also been theoretically validated. Based on the experimental results on a series of synthetic and real data sets, it has been substantiated that the proposed algorithms attain very robust performance on data with outliers and missing components. As compared with the current state-of-the-art methods, our algorithms exhibit notable advantages in both computational speed and accuracy. The experimental results also illuminate the potential usefulness of the proposed algorithms on large-scale face recognition and SFM applications.

REFERENCES

- [1] M.J. Black and A.D. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. *International Journal Computer Vision*, 26(1):63-84, 1998.

- [2] A.M. Buchanan and A.W. Fitzgibbon. Damped Newton algorithms for matrix factorization with missing data. CVPR, 2005.
- [3] F. De La Torre and M.J. Black. A framework for robust subspace learning. International Journal Computer Vision, 54(1-3): 117-142, 2003.
- [4] C. Ding, D. Zhou, X.F. He, H.Y. Zha. R1-PCA: Rotational invariant L1-norm principal component analysis for robust subspace factorization. ICML, 2006.
- [5] A. Eriksson and A.v.d. Hengel. Efficient computation of robust low-rank matrix approximations in the presence of missing data using the L_1 norm. CVPR, 2010.
- [6] A. Georghiades, P. Belhumeur, D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE Trans. PAMI, 23(6): 643-660, 2001.
- [7] G. Golub and C.V. Loan. Matrix Computation. Maryland: Johns Hopkins University Press, 1989.
- [8] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. Journal of the Optical Society of America A, 11(11): 3079-3089, 1992.
- [9] Q. Ke and T. Kanade. Robust L_1 norm factorization in the presence of outliers and missing data by alternative convex programming. CVPR, 2005.
- [10] D.E. Knuth. The Art of Computer Programming, Volume 3: Sorting and Searching. Massachusetts: Addison-Wesley, 1973.
- [11] N. Kwak. Principal component analysis based on L1-norm maximization. IEEE Trans. PAMI, 30(9): 1672-1680, 2008.
- [12] Z.C. Lin, M.M. Chen, Y. Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-Rank matrix. Technical Report UILU-ENG-09-2215, 2009.
- [13] T. Okatani and K. Deguchi. On the Wiberg algorithm for matrix factorization in the presence of missing components. International Journal Computer Vision, 72(3): 329-337, 2007.
- [14] N. Srebro and T. Jaakkola. Weighted low-rank approximations. ICML, 2003.
- [15] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. International Journal Computer Vision, 9(2): 137-154, 2007.
- [16] J. Wright, Y.G. Peng, Y. Ma, A. Ganesh, S. Rao. Robust principal component analysis: Exact recovery of corrupted low-rank matrices by convex optimization. NIPS, 2009.